

AN ARCHITECT OF THE GOLDEN YEARS

MICHAEL D. ZEILER

EMORY UNIVERSITY

William H. Morse has played a major role in the experimental analysis of behavior. His view of operant behavior as the outcome of differential reinforcement provides an invaluable lesson in scientific research and theory. He studied schedules of reinforcement to generate an in-depth analysis of the complex interactions existing when contingencies exert their control over behavior. He has been instrumental in showing how behavior is determined by the dynamic interaction of factors brought into play by the imposition of any schedule, and he has a remarkably intuitive understanding of the nature of these determining variables. Some of these causal events are imposed directly by the schedule, but others arise in a more indirect manner through necessary constraints. In Morse's view, schedules can be more fundamental in determining behavior than are the scheduled events themselves. Behavior is the shaped product of an organism's history in combination with present environmental conditions. His impact deserves to be more than historical: A study of his work continues to reward the reader with exciting insights into the nature of behavioral control.

Key words: William H. Morse, schedules of reinforcement, schedule theory, operant behavior, shaping

The 1950s and 1960s were an exciting period in the experimental and theoretical analysis of behavior. After years of domination by experimental psychologists devoted to testing global learning theories, a small group then deemed "Skinnerians" or "operant conditioners" set a new path. Unencumbered by traditional learning theory, they developed new techniques that led to extraordinary advances in understanding how behavior was produced and maintained. Starting with the first issue of the *Journal of the Experimental Analysis of Behavior (JEAB)* in 1958, every succeeding issue in the next 10–15 years described interesting and exciting advances that made readers eager for the next one. Whether the topic was schedules of reinforcement, conditioned reinforcement, or stimulus control, the advances made during this period were major breakthroughs and testimonials to the power of the experimental analysis of behavior in clarifying and exposing the principles of behavior. William H. Morse was a leading contributor to this achievement.

Morse received his doctorate in Psychology at Harvard where his advisor was B. F. Skinner. This transplanted Virginian then chose to remain in the Boston area when he took a position in the Department of Pharmacology at Harvard Medical School and became a col-

league of Peter Dews. Roger Kelleher joined them a few years later. The impact of this remarkable threesome is highlighted in this issue of *JEAB*; my purpose here is to discuss and emphasize Morse's important contributions.

Volume 1 (1958) of *JEAB* offers vivid examples of his experimental style. (Morse was on the first Board of Editors and later became an Associate Editor of the journal.) I should preface my discussion of these papers by mentioning that Morse was known to be an unusually skilled shaper of behavior. He would spend hours watching experimental animals and in developing control by the contingencies of interest. Only a few operant conditioners were acknowledged to be skilled architects of behavior, and Morse was a member of this select group who seemed able to intuit what was required to generate and maintain behavior. The knack seems to be a talent that you either do or do not have; I have never heard of anyone who was taught this kind of exquisite sensitivity.

One early experiment posed the question of how long an animal could be kept responding in a continuous session (Skinner & Morse, 1958b). The chosen schedule was a differential-reinforcement-of-low-rate (DRL), because it would sustain responding even with very infrequent reinforcer deliveries. Whenever the pigeon let 1 minute pass without pecking the key, the next peck would result in a food presentation (DRL 1-min schedule). Under

Correspondence may be sent to: Michael D. Zeiler, Department of Psychology, Emory University, Atlanta, GA 30329 (e-mail: psymdz@emory.edu).
doi: 10.1901/jeab.2006.24-06

still longer schedules (DRL 3–4 min), the pigeons rarely met the pausing requirement. Instead, they responded continually, and thereby received very little food. One pigeon's rate was such that it experienced 31 *hours* between food deliveries, never allowing the required time to pass without a response. So, the point was made that sessions could be very long indeed with extremely low rates of reinforcer delivery (and that perhaps pigeons never need to sleep).

One aspect of their procedure merits comment. Why did the experimenters choose to program 5.5 hours of DRL separated by fixed-interval (FI) and fixed-ratio (FR) schedules that lasted for just one food delivery each (thus making the arrangement a multiple schedule)? The official story is that they also were interested in whether such schedules would have their characteristic effects in a continuous session. My suspicion is that Morse somehow intuited that this was the right thing to do in order to maintain any significant responding with the DRL schedule. Otherwise, why bother with the interposed schedules?

Another experiment involved fixed-interval reinforcement of running in a wheel (Skinner & Morse, 1958a). After noting discrepancies between the temporal patterning of the running response displayed by rats and those patterns observed with bar-pressing, Morse played with the friction against which the wheel turned. He found that with the right friction the fixed-interval pattern looked identical to bar-pressing. What led to that manipulation? Why should friction have unlocked the effect? Did Michelangelo or Picasso know exactly why they did everything they did that made their work so unique and powerful? Can it be learned or taught? To some extent, perhaps. For example, I once asked Morse about the proper adjustment for a pigeon key. His answer was to suggest that I set up the key in a comfortable position for myself and then adjust the force setting until it felt right to me when pressing it as fast as I could. That setting would be the right one for a pigeon as well. So now I knew how a master calibrated a key, but to this day I still didn't know why this was the right thing to do or how he arrived at this insight.

Most of us probably do not view the control of behavior as an art form, but how else should we view innovative decisions made about what to do when trying to develop or explicate

a novel behavior? A number of Morse's papers refer to using the right parameter values in developing behavior. What makes a parameter value right? I suspect most researchers view parameters as variables having continuous effects on behavior, whether those effects be linear, non-monotonic, or the like. The situation is quite different, however, if parametric variations introduce new controlling variables and processes. The implication is that distinctive effects derive from the interactions of certain key parameters, and the wrong parameters will lead to erroneous conclusions about the general effects of controlling variables. Behavior may occur with one parameter value but not with another. It was never clear, however, how one was to decide what values to use or how different factors would interact with each other at different values. The approach seemed largely intuitive, and the proof of the pudding for those skilled in developing and maintaining behavior was in the appearance and maintenance of the behavior of interest.

In another example, Marley and Morse (1966) used newly hatched chickens as subjects. Chicks were used to study the behavioral effects of certain drugs because the blood-brain barrier had not yet developed in these young animals. Not much was known about operant conditioning in baby chicks, so it was necessary to develop appropriate procedures. A serious problem emerged. When socially isolated as they must be to examine how schedules and drugs affect their behavior, newly-hatched chicks emit sustained distress cheeping and seem too depressed to do anything but cry and hide. Conditioning seemed impossible until Morse put a mirror in the experimental chamber. This addition made the previously depressed chicks apparently happy. Alleviating depression with a mirror may not testify to the intellectual power of chickens, but it did eliminate distress cheeping and allowed the chicks' behavior to be conditioned. How did Morse happen to think of that? All he says in that regard is that it worked, but then he did go on to follow up on the finding by exploring in detail the effects of mirror removal and replacement.

An extraordinarily counterintuitive finding emerging from the continuously productive collaboration of William Morse and Roger Kelleher was that behavior could be main-

tained by response-produced electric shock. I can't remember whether it was Morse or Kelleher who told me that the first procedure was intended to generate a pattern of negatively accelerated responding that would have been useful in testing their rate-dependency hypothesis of how drugs affect behavior under many schedules of reinforcement. The procedure involved a variable-interval 3-min (VI 3-min) schedule of food delivery imposed conjointly with a fixed-interval 10-min (FI 10-min) schedule of shock presentation. Each response during minutes 10–11 produced a strong shock, but all shocks could be avoided if the animal withheld the response during that minute. In schedule terminology, the shock schedule was FI 10-min of shock presentation with a limited hold of 1 min imposed conjointly with a VI 3-min schedule of food delivery. The surprising result was that negatively accelerated responding did not occur. Instead, responding was positively accelerated until the first response after 10 min produced the shock, and then responding stopped for the next minute. This surprising finding apparently caused Kelleher and Morse to lose interest, at least temporarily, in using shock schedules to test rate-dependency and led instead to intense scrutiny of the conditions allowing behavior to be maintained by response-produced electric shock. The procedure and results now are part of common behavior-analytic knowledge (Kelleher & Morse, 1968; Morse & Kelleher, 1977).

They went on to show that the reinforcing ability of strong electric shock with fixed-interval schedules required no history of food reinforcement. It also would occur if the history first involved shock postponement, or simply shock presentation, or other procedures as well. Nor was it the case that the phenomenon derived from "fooling" the animal about shock. Morse spent hours watching the monkeys and tailoring procedures based on his intuitions of what would make shock function as a positive reinforcer. These demonstrations each were followed by detailed experiments that proved his intuitions correct. Behavior maintained by response-produced electric shock was an important component leading to the conclusion that schedules themselves are fundamental determinants of behavior, seemingly more important than the particular events scheduled (cf. Morse &

Kelleher, 1977). All reinforcers, even the familiar ones of food, water, and sex, derive from history and ongoing behavior; there is nothing special about electric shock in that regard. Implicit in these findings is the need to reconceptualize reinforcing events in terms of schedule effects. A focus on the dynamics of acquisition and maintenance of characteristic patterns of responding under a schedule is a far more sophisticated perspective on the meaning of reinforcement than simply increases in response probability.

In addition to Morse's ability as a working analyst of ongoing environment-behavior interactions and their history, and his ability to shape behavior, he also was a first-rate theorist. The first issue of *JEAB* includes a paper on conjunctive fixed-interval fixed-ratio schedules of food presentation (Herrnstein & Morse, 1958). Under this conjunctive schedule, food is presented when responding has met both the FI and the FR requirements. For example, if the schedule is conjunctive FI 15-min FR 40, the animal receives food for the first response occurring after 15 min as long as it has emitted at least 39 responses during that interval. If it has not responded enough during the fixed interval, it receives food as soon as it emits the 40th response and thereby completes the fixed ratio. When the authors systematically manipulated the size of the ratio requirement, they found that the highest response rate occurred with a simple FI 15-min schedule, and then declined with each successively larger fixed ratio (FR 10, 40, 120, and 240). The theoretical explanation was that the ability of fixed-interval schedules to maintain large amounts of unreinforced behavior derives from the fact that it places no constraint on response output prior to food delivery. However rare, reinforcement following intervals containing few responses appears essential for the FI schedule to maintain a high overall number of responses.

Here is theory closely tied to data. This approach clearly characterizes the theory that schedules are fundamental determinants of behavior. But it achieves its most sophisticated level in Morse's (1966) penetrating and challenging theory of how schedules come to exert their powerful effects on behavior. In my opinion, this chapter is one of the most important papers in the history of learning theory as well as in the history of the

experimental analysis of behavior. For the former, it is a lesson in how learning theory can be conducted effectively and productively with a minimum of unfounded speculation and a maximum of clear thinking; for the latter, it is a powerful testimony to a breadth and depth of analysis that unfortunately no longer seems to have the central importance to the field it once did. Much of the remainder of this paper will focus on this chapter, its contributions and questions it raised, many of which remain unresolved and await further research (see also, Zeiler, 1977, 1979).

Morse (1966) viewed schedule performances as the products of shaping. "Conditioned operant behavior emerges out of undifferentiated behavior through successive approximations to new and more complex forms by the process of successive differential reinforcement (shaping). Behavior that has become highly differentiated can be understood and accounted for only in terms of the history of reinforcement of that behavior—when, and how, and under what stimulus conditions reinforcers acted to shape the behavior" (p. 52). The effect of a schedule of reinforcement is to establish a highly predictable temporal pattern of responding, and this sequential organization inevitably appears at the right parameter values despite differences in responses and reinforcers and in whether overall response rate increases or decreases. (For example, in the experiment mentioned above on fixed-interval reinforcement of running in a wheel, running showed the characteristic fixed-interval pattern, even though the overall rate of running decreased from its baseline level.) The central theoretical question is how such predictable sequences of behavior are generated automatically simply by imposing the schedule of reinforcement. As Morse (p. 61) puts it, a central issue is to determine "how a schedule of intermittent reinforcement operates upon behavior to engender characteristic patterns of responding." The focus is on the contingencies operating to produce the totally predictable effects of ratio and interval schedules.

A given schedule has its effects because of the relations between responding and reinforcement which inevitably follow simply because that schedule has been imposed. For example, FR schedules specify constancy in the number of responses and that the time to

reinforcement depends on response rate, whereas FI schedules specify the minimum time to reinforcement and render high response rates essentially irrelevant. The lack of specification of time between reinforcers on FR and number of responses per reinforcer on FI may in itself be as important as the required events involved in the schedule. In addition, what the animal is doing at the moment of reinforcement may vary in successive ratios or intervals, and these features also play a role. Periods of pausing on interval schedules are often followed by reinforcer presentation after a small number of responses, but this rarely if ever happens with ratio schedules. As Morse knew well from his early work on superstitious discriminations (Morse & Skinner, 1957), the conditions prevailing at the time of reinforcement gain control over behavior whether or not the experimenter has specified those conditions. The uniformity of schedule-controlled behavior within and between subjects shows that the sequential interactions between the schedule and responding must be inevitable despite different starting points. As Morse (p. 77) says, "A simple schedule is one that is simple to specify and program rather than one that has a simple relation to behavior."

Morse focuses on reinforcers operating at different levels of behavior. Reinforcers not only influence the response that precedes them, but also may affect the entire preceding sequence of behavior. Behavior also may be sensitive to overall densities or probabilities of reinforcement. The understanding of how schedules exert their particular effects entails consideration of these multiple sources and levels.

The first feature considered is the reinforcement of different interresponse times (IRTs). Neither ratio nor interval schedules specify the IRT that must precede reinforcement. Skinner (1938) had suggested that different IRTs are likely to be reinforced with ratio and interval schedules. With ratio schedules, a response with any IRT is equally likely to be reinforced, but in interval schedules, the longer the IRT the more likely it is to be reinforced. Morse (1966) adds the interesting comment (Note 6, p. 70) that an FI schedule can be defined as a schedule reinforcing a minimum cumulative sum of IRTs, and reminds us that this definition reduces concern with why so many responses occur on these schedules when only

one is required. It also may provide some insight into why responses occur at a fairly high rate throughout the interval.

Is differential reinforcement of IRTs a plausible candidate for explaining rate differences in different schedules? The starting point is that explicit reinforcement of longer IRTs in DRL schedules clearly affects response rate. Furthermore, requiring the animal to meet a DRL requirement at the end of an FI, VI, or FR schedule reduces overall response rate. Such data prove that IRT reinforcement can be important. However, showing that explicit IRT reinforcement influences responding does not mean that IRT reinforcement is always exerting a significant effect in the absence of a specific IRT requirement. One must look at patterning as well as at response rate. Appropriate schedule patterning is maintained only with VI schedules when IRT requirements are added. With FI and FR schedules, explicit IRT requirements are likely to change the pattern from that which occurs in their absence. Given those data, it might be reasonable to conclude that normal FI and FR performances are not strongly affected by differential reinforcement of specific terminal IRTs. When the terminal IRT must be shorter than a specified value (DRH schedules), response rates often increase, but typically a number of unusual and unexpected effects are also generated. More reliable rate-enhancing effects are seen when a sequence of responses must be emitted in less than a specified time than when the single terminal response must conform to a DRH requirement. These observations not only highlight the complexity of schedule interactions but show that it is possible to disentangle the sources of control. Although shorter IRTs are reinforced on FR than on FI schedules, it is not clear that this is more than the necessary outcome of terminal response rates being higher with ratio schedules. The reason for these rate differences probably should be sought elsewhere.

Where IRT reinforcement considers the properties of the terminal response in the sequence of behavior leading to reinforcer delivery, it also is possible to consider reinforcement in terms of averages. The role of average frequency of reinforcement in determining response rate is far more complicated than is acknowledged by contemporary

theory, whether the focus is on behavior averaged over sessions or groups of sessions, or on the behavior leading to individual reinforcer presentations. Interval schedules allow the direct manipulation of the time between successive reinforcers. Averages of these times suggest that response rates increase with higher reinforcer frequencies in both VI and FI schedules. If, however, a small FR requirement is added after an FI schedule has been completed (tandem FI FR schedule), average response rates increase even though this schedule reduces reinforcer frequency from what it is when an FI schedule is imposed alone. Further, although reduced reinforcer frequency produces consistently lower average rates on VI schedules, it does not have the same effects with FR schedules. With ratio schedules, a lower reinforcer frequency first results in rate increases and then finally in rate decreases, probably because other factors become increasingly important. The average frequency of reinforcement in time alone does not explain average response rate in all schedules, although it may play a role in some.

The idea that reinforcer frequency is a ubiquitous controlling variable is even more seriously challenged as soon as attention shifts from average response rate to the behavior leading up to each reinforcer presentation. It is hardly news that averages taken over many instances of behavior may grossly misrepresent the individual instances of the behavior that is being averaged, but perhaps not everyone is aware of how vividly schedule performances emphasize that point. In FI schedules, average response rate decreases with longer values (decreases in reinforcer frequency). Now consider the performance in individual intervals. Intervals with high rates commonly are followed by intervals with low rates even though reinforcer frequency is unchanged from one to the next (cf. Skinner, 1938). Something other than frequency clearly is at work in determining ongoing behavior; indeed, this something seems to have the ability to overwhelm any effects that the constant interreinforcer times might have on response rate. Could the reason for the effects of interval schedules on response rate depend on the fact that they let the number of responses vary from one interval to the next? If so, why should that be the case? And what does that mean for FR schedules where such

variability is precluded? Are VR schedules more like FI schedules than they are like FR because they, too, result in variability in the number of responses emitted per reinforcer, or is it important whether or not the schedule forces variability or leaves it up to the animal? Taking another case, in VI schedules long intervals may be followed by an increased response rate, reminiscent of the early effects of extinction. What does this imply about the relation between average reinforcer frequency and average response rate, which seems to assume that VI schedules produce stable response rates from one interval to the next?

It is perhaps not surprising for a researcher who spent hours watching experimental animals and shaping their behavior to emphasize the properties of single reinforcers in generating schedule-controlled behavior. Even those less skilled than Morse in shaping behavior can testify to the immediacy of the effects of differential reinforcement. Furthermore, watchful students of steady-state behavior under schedules of reinforcement know that transitions occur rapidly when schedules are changed. Changing the size of a fixed interval results in changes in behavior after one or two reinforcers, and changing from an interval to a ratio schedule generates virtually instantaneous increases in response rate. When a very large fixed ratio has essentially reduced responding to a rare event, one reinforcer given on an FI schedule regenerates behavior almost immediately. Such rapid effects do not argue in favor of hypotheses that steady-state behavior is the product of molar variables that operate over many reinforcers; they argue for ongoing dynamic effects of individual reinforcers. Morse seems always to have known that individual reinforcers operate in an ongoing and selective manner to generate and maintain all operant behavior.

It should not be surprising, therefore, that Morse emphasizes momentary effects of reinforcers even though he considers the possibility of control by more molar determinants. When Morse talks about reinforcer frequency, his primary reference is to the just-experienced IRT; when he talks about responses per reinforcer, he is referring to the number of responses that just occurred; when he treats the probability of an IRT being reinforced, it is in the context of the distribution of the IRTs that were reinforced rather than in terms of

the average reinforced IRT. To a considerable extent, contemporary theories of operant behavior tend to ignore individual reinforcers in favor of averages. This probably has followed from the success of various forms of the matching law in describing choice in concurrent schedules and average response rate in VI schedules. It should be noted that there is far less evidence for a systematic effect of VI size on response rate when each schedule is maintained alone until response rate stabilizes than when the same subject is exposed to many different average values in the same session. In the first case, the data reported sometimes show no effect of different VI schedules and sometimes show disorder; in the second case, response rate systematically varies with the size of the VI. Different dynamic variables interact to determine behavior occurring when an animal is exposed to two or more schedules at a time than when it has single schedules in effect alone. Unfortunately, no one yet has explained the differences. I would recommend that anyone interested in pursuing this issue consult Morse's (1966) chapter for some interesting suggestions about what is probably going on.

The student of environmental contingencies who uses schedules as the vehicle for understanding them surely cannot deal only with response rate, because schedules also differ so noticeably in the patterns of responding they generate. It has been tempting for some to attribute the different patterns to the presence or absence of stimulus control by elapsed time to reinforcement, that is, to the discrimination of time intervals. However, when IRT requirements added to FI and FR schedules disrupt the normal temporal distributions of responding without markedly changing the time to reinforcement, it seems evident that more than temporal discrimination is involved in patterning. In fact, Morse downplayed the role of temporal discrimination in determining the distribution of responses in time. "Our understanding of schedule performances is fettered by the tendency of many authors to ignore the dynamic equilibrium conditions inherent in responding under a schedule and to explain schedule performances instead as discriminations of subtle differences in inferred stimulus conditions (Morse, 1966, p. 86)." If temporal discrimination does come into play in some schedules, it always is in

conjunction with other and probably more important factors. The characteristic temporal distributions of responses follow from the dynamic interaction of multiple sources of control such as reinforcer frequency, the nature of reinforced IRTs, the need to respond in order for the reinforcer to occur, delayed reinforcement for responses prior to the last, with the possibility that stimulus duration also plays some role. As with response rate, patterning is the outcome of multiple subtle variables.

Attributing behavior to averages of some property of reinforcement or to any single factor just cannot do justice to the complexity of any schedule effect. Averages clearly do not explain immediate dynamic effects, they do not help us understand the quickness of transition states, and, most importantly, they simply do not describe behavior as it occurs in real time. Therefore, they do not explain how schedules exert their effects on behavior. What also needs to be recognized is that such molar rules as matching response allocations to reinforcer allocations or maximizing average reinforcer frequency are the outcomes of behavior, not its antecedent cause. If we were to find that these rules meant that animals always respond in a way that maximizes return while minimizing energy output, we would understand why those outcomes would have been favored by natural selection. But we still would not know the momentary events that guide the behavior as it occurs. That is the challenge that Morse poses and tries to meet. He begins with his intuitions and then goes on to analyze the immediate events that control behavior, and that is why he was a shaper and behavioral theorist *par excellence*.

Morse's approach to behavior has far more than historical significance. What he did and said was and still is important because we have not resolved the problems he poses so well for us. To understand contingencies is to know how they exert their effects. Do we know more about shaping now than we did in 1970? What about our knowledge of how schedules exert their effects? What about the dynamics of stimulus control? These and other core areas still pose basic and unanswered questions that can be answered with clear thinking and good data. Maybe if we and our students remind ourselves of what we expected from the study of behavior by rereading and studying Morse's

work, a new and exciting flame could be rekindled. This is exactly what happened to me in the course of writing this essay!

I also wish to comment on William H. Morse as a person. I never met anyone who worked in the Harvard Medical School laboratory who was not indebted to him as a teacher and scholar and model of a simultaneously rigorous and insightful researcher. Add to that the fact that everyone knows him as a warm and caring person who has done some extraordinary things for other people. My own experience began as a callow new Ph.D. teaching at Wellesley College who was invited to attend the Friday laboratory meetings at the Medical School and had the privilege of being tutored in operant conditioning by Bill and the other giants in that laboratory in the most positive and supportive way possible. For me, as for the many who came before and after, the experience was unforgettable.

REFERENCES

- Herrnstein, R. J., & Morse, W. H. (1958). A conjunctive schedule of reinforcement. *Journal of the Experimental Analysis of Behavior*, 1, 15-24.
- Kelleher, R. T., & Morse, W. H. (1968). Schedules using noxious stimuli. III. Responding maintained with response-produced electric shocks. *Journal of the Experimental Analysis of Behavior*, 11, 819-838.
- Marley, E., & Morse, W. H. (1966). Operant conditioning in the newly hatched chicken. *Journal of the Experimental Analysis of Behavior*, 9, 95-103.
- Morse, W. H. (1966). Intermittent reinforcement. In W. K. Honig (Ed.), *Operant behavior: Areas of research and application* (pp. 52-108). New York: Appleton-Century-Crofts.
- Morse, W. H., & Kelleher, R. T. (1977). Determinants of reinforcement and punishment. In W. K. Honig, & J. E. R. Staddon (Eds.), *Handbook of operant behavior* (pp. 174-200). Englewood Cliffs: Prentice-Hall.
- Morse, W. H., & Skinner, B. F. (1957). A second type of superstition in the pigeons. *American Journal of Psychology*, 70, 308-311.
- Skinner, B. F. (1938). *The behavior of organisms*. New York: Appleton-Century-Crofts.
- Skinner, B. F., & Morse, W. H. (1958a). Fixed-interval reinforcement of running in a wheel. *Journal of the Experimental Analysis of Behavior*, 1, 371-379.
- Skinner, B. F., & Morse, W. H. (1958b). Sustained performance during very long experimental sessions. *Journal of the Experimental Analysis of Behavior*, 1, 235-244.
- Zeiler, M. D. (1977). Schedules of reinforcement: The controlling variables. In W. K. Honig, & J. E. R. Staddon (Eds.), *Handbook of operant behavior* (pp. 201-232). Englewood Cliffs, NJ: Prentice-Hall.
- Zeiler, M. D. (1979). Output dynamics. In M. D. Zeiler, & P. Harzem (Eds.), *Reinforcement and the organization of behavior* (pp. 79-115). London: Wiley.